

Why Your Data Is At Risk

by Randy Nash - @RISK Online - Monday, 20 December 2004.

Electronic data resides in two basic areas:

- In bulk in some form of repository, such as a database or collections of individual files (called data at rest)
- In small quantities being transmitted over a network (called data on the wire)

Your data is vulnerable no matter where it resides. While most companies take security precautions, many of those precautions turn out to be insufficient to protect valuable corporate assets. The key lies in knowing where vulnerabilities exist and making appropriate risk-based decisions.

Introduction

The ability to gather and share volumes of information was the primary reason behind the creation of the Internet, but such wide availability greatly magnifies the risk of that information being compromised. Attacks against large databases of critical information are on the rise, such as in the following recent cases:

- February, 2003: A hacker broke into the security system of a company that processes credit card transactions, giving the hacker access to the records of millions of Visa and MasterCard accounts.
- June, 2004: More than 145,000 blood donors were warned that they could be at risk for identity theft from a stolen university laptop containing their personal information.
- October, 2004: A hacker accessed names and social security numbers of about 1.4 million Californians after breaking into a University of California, Berkeley computer.

NOTE - *Identity theft* occurs when someone uses your personal information—such as your name, social security number, credit card number, or other identifying information—without your permission, frequently to commit fraud or other crimes.

Vulnerabilities of Data on the Wire

Data on the wire is vulnerable to some very focused attacks. Data can be intercepted (sniffed). ARP attacks can be used to sniff information in a switched environment. ARP attacks can also be used to initiate "man in the middle" attacks that can allow an attacker to intercept and potentially modify information in transit.

Sniffing

Sniffing refers to a technique for capturing network traffic. While sniffing can be accomplished on both routed and switched networks, it's much easier in a routed environment:

- Layer 3 devices, such as routers, send information by broadcasting it to every destination on the network, and the destination handles the problem of parsing out the specific information that's needed from the general broadcast.
- In a switched environment, switches send traffic only to

its intended host (determined by the destination information in each individual packet).

Operating in a switched environment doesn't totally alleviate the risk of sniffing, but it does mitigate that risk to a large degree.

Most networks today also utilize virtual LAN (VLAN) configurations to segment network traffic and further reduce the risk of sniffing. A VLAN is a switched network that's logically segmented. VLANs are created to provide the segmentation services traditionally provided by routers in LAN configurations. VLANs address scalability, security, and network management. Routers in VLAN topologies provide broadcast filtering, security, address summarization, and traffic-flow management.

Just as switches isolate collision domains for attached hosts and only forward appropriate traffic out a particular port, VLANs provide complete isolation between VLANs. None of the switches within the defined group will bridge any frames—not even broadcast frames—between two VLANs. Thus, communication between VLANs is accomplished through routing, and the traditional security and filtering functions of the router can be used.

Segmentation can be organized in any manner: function, project team, application. This capability is especially useful for isolating network segments for security purposes. For example, you may place application servers on one VLAN and system administrators on another (management-level) VLAN, with access control lists to restrict administrative access to only that VLAN. This setup can be accomplished regardless of physical connections to the network or the fact that some users might be intermingled with other teams.

ARP Attacks

The Ethernet Address Resolution Protocol (ARP) enables systems to find the unique identifier (MAC address) of a destination machine. ARP attacks provide the means to either break or misuse the protocol, with the goal of redirecting traffic from its intended destination. In an ARP attack, the attacker can sniff, intercept, and even modify traffic on a compromised network segment.

The effectiveness of these attacks is limited in two ways:

- **Data on the wire is generally available only in small pieces.** It's true that many systems and applications send login/password pairs in clear text (without any encryption). An attack may capture such small bits of data; it may even be possible over time to assemble enough useful information to make identity theft possible. However, the attacker must either be directly connected to the internal network, or have succeeded in compromising an internal system and installing some form of sniffer to gather information. For the effort to be worthwhile to the hacker, many small chunks would need to be captured and then filtered out of the massive volumes of traffic traversing most of today's networks; and then the captured data would have to be reassembled into meaningful information. This is a tremendous task with a potentially very small payoff.
- **Capturing data takes time.** The longer the attacker is

Capturing data takes time. The longer the attacker is inside the network, the more likely he or she is to get caught. It's easier to get information at the source, rather than trying to capture and decode thousands of network packets.

Vulnerabilities of Data at Rest

While sniffing data on the wire may yield a big reward, data at rest is the proverbial pot of gold. Most organizations maintain detailed databases of their personnel information, for example, making the large corporation a very tempting target. These databases regularly contain quantities of names, addresses, and even social security numbers for tax purposes. This is all the information that someone needs to steal your identity. Statistics show that identity theft attacks are increasing. More than thirty thousand victims reported ID theft in 2000; in 2003, the Federal Trade Commission received more than half a million complaints.

A major issue in protecting your data repository is the fact that there are so many avenues of attack. Attacks can be launched against the operating system, the database server application, the custom application interface, the client interface, and so on. Application attacks don't have to be directed against the target application, either. Any attack providing system-level access to an attacker is a risk to your data.

Your system is also a potential target for a multitude of computer viruses, worms, and Trojans. Current reports put the number of these types of applications at more than 100,000. Many recent computer worms leave systems vulnerable by covertly installing a backdoor that enables the attacker to enter the system at will.

How Can We Protect Our Data?

How do we defend against so many possible attack vectors? The key is to focus on the data. The first step should be data-sensitivity analysis as part of an overall risk-assessment process. Data-sensitivity analysis begins by identifying an organization's critical data and ways in which that data is used. Once the sensitivity of data has been classified, the organization can reach decisions about the necessary level of protection for that data. Your tendency may be to apply the greatest level of protection available, but that level may be neither necessary nor cost-effective. For example, you wouldn't spend \$100,000 on a firewall to protect an expected loss of only \$5,000. You can get a better idea of how to apply countermeasures if you include a loss/impact analysis as part of the risk-assessment process.

Simple Approach

A simple approach to data protection looks at the various layers of security that can be applied. Consider the following starting checklist:

Data repository:

- Do you need to encrypt the data repository?
- Do you need a hash of the transactions for integrity purposes?
- Should you digitally sign transactions?
- Make sure that database logging is enabled and

properly configured.

Server considerations:

- Harden the operating system.
- Disable unnecessary services and close ports.
- Change system defaults.
- Don't use group or shared account passwords.
- Lock down file shares.
- Restrict access to only necessary personnel.
- Consider host-based firewalls and intrusion detection for critical servers.
- Maintain proper patch procedures.

Network segment:

- Use switches rather than routers or hubs as much as possible.
- Lock down unused router/switch ports.
- Consider MAC filters for critical systems.
- Establish logical subnets and VLANs.
- Set up access control lists (ACLs) for access routes.
- Use ingress/egress filters, anti-spoof rules.
- Determine appropriate location and functionality for network-based firewalls and intrusion detection.
- Use encrypted logins or SSL for web-based sessions.

Physical security for data:

- Establish input/output handling procedures.
- Use physical access logs for server rooms and network operations centers.
- Document tape-handling procedures, tape rotation, offsite storage.
- Consider an alternate data center.
- Archiving: Where does your data go to rest in peace?
- Data destruction: Degauss, erase/overwrite, physical destruction?
- How is data handled when equipment is sent out for repair, replacement, or end of life?

This is just a quick list of points to consider. Fortunately, folks much smarter than I am have developed a much more comprehensive approach.

Structured Approach

Security standards and guidance are available, especially for organizations that are part of or do business with the U.S. government. Through the work of various organizations, the government has put together a program known as Certification & Accreditation (C&A). Standards have been and continue to be developed that provide guidance on the performance of risk assessments, development of security plans, and the application of security controls.

The Computer Security Division of the National Institute of Standards and Technology (NIST) has been assigned this important multi-part task:

- Improving federal information-systems security by raising awareness of IT risks, vulnerabilities, and protection requirements, particularly for new and emerging technologies.
- Researching, studying, and advising agencies of IT vulnerabilities and devising techniques for the

cost-effective security and privacy of sensitive federal systems.

- Developing standards, metrics, tests, and validation programs.
- Developing guidance to increase secure IT planning, implementation, management, and operation.

The C&A process is explained and documented in NIST's [publications](#). NIST's guidelines provide an excellent framework for selecting, specifying, employing, and evaluating the security controls in information systems.

Summary

Data is under constant attack from a growing number of sources. It's vital that you know what data you have, how sensitive that data is, how critical it is to your corporate mission, and the risks it faces. Perform a risk assessment, and, once the threat level has been determined, develop an appropriate plan to protect that data with multiple layers of security. Only by being aware of your valuable assets can you properly monitor and protect them.

